98-4727

Issues in Industrial Multidisciplinary Optimization

J. Bennett, P. Fenyes, W. Haering, M. Neal
Body Engineering and Integration Department
General Motors Research and Development Center
30500 Mound Road Warren, MI 48090-9055

Abstract

Several mathematically based multidisciplinary design strategies are illustrated with an exploratory multidisciplinary analysis and optimization package on a simple example problem. These examples are used to motivate a discussion of required data handling and processing modules. These requirements envision a situation where some disciplines may have computationally expensive analysis capabilities and where not all disciplines have easily available approximations for all required quantities

Introduction

There are generally two areas of development in multidisciplinary optimization and design systems. The first is the formal mathematical approach that is generally characterized by the work presented at the Multidisciplinary Design and Optimization Conferences. The second is a more ad hoc approach which is evolving from the traditional design and analysis communities and is typified more by a Multidisciplinary Analysis capability that is evolving in the commercial CAD and CAE environments. These environments envision a common parametric description of the artifact and an ability to generate input information for several disciplines from this format. Then analyses will be performed using complex commercial or proprietary codes and decisions made on how to modify the initial design. This process is usually characterized by significant human interaction to develop the artifact model, generate the analysis models, execute the analysis models and finally to examine the output and make decisions. The formal mathematical approach tends to use much simpler and easily modified local analysis methods (that execute on the order of minutes or seconds) and to concentrate on multidisciplinary design algorithms which interact with the analysis methods in an almost automatic fashion.

This paper will discuss some of the issues associated with developing an industrial, rapid, multidisciplinary design system that makes use of some aspects of modern multidisciplinary optimization research while being constrained by analysis software packages that do not all have consistent local optimization capabilities and are not easily modified.

The work is conducted in the experimental Integrated Vehicle Design Analysis (IVDA) system that has been developed at GM R&D Center over the past few years. This system was described in some detail in [1]. Only that detail which is critical to the present discussion will be included here. A flow chart of the complete system is shown in Figure 1. This system envisions a parametric description format that for a specific instantiation of the parameters will generate a common vehicle description that in turn is used to generate input for an extensive set of disciplinary analysis capabilities. Note that this vehicle database contains more than just a geometric description of the vehicle in that it includes materials and their properties as well as mass and inertia characteristics of various components. The geometric design parameters include both global vehicle dimension and component structural dimensions as shown in Figure 2. The various analysis tools represent a range of capabilities. There are both commercial and proprietary codes and some disciplines have design (optimization) capabilities and others do not. The analysis capabilities in each discipline were selected to represent a preliminary analysis capability. In most cases they represent neither the simplest nor the most complex analysis capability in each discipline. They do characterize the current state of computer based engineering analysis. An initial goal of the system was to be able to complete one full analysis cycle in 24 hrs. For many of the disciplines the development of the input data is considered to be the major time constraint. For this reason, highly automated model generation methods based on templates were developed in each discipline.

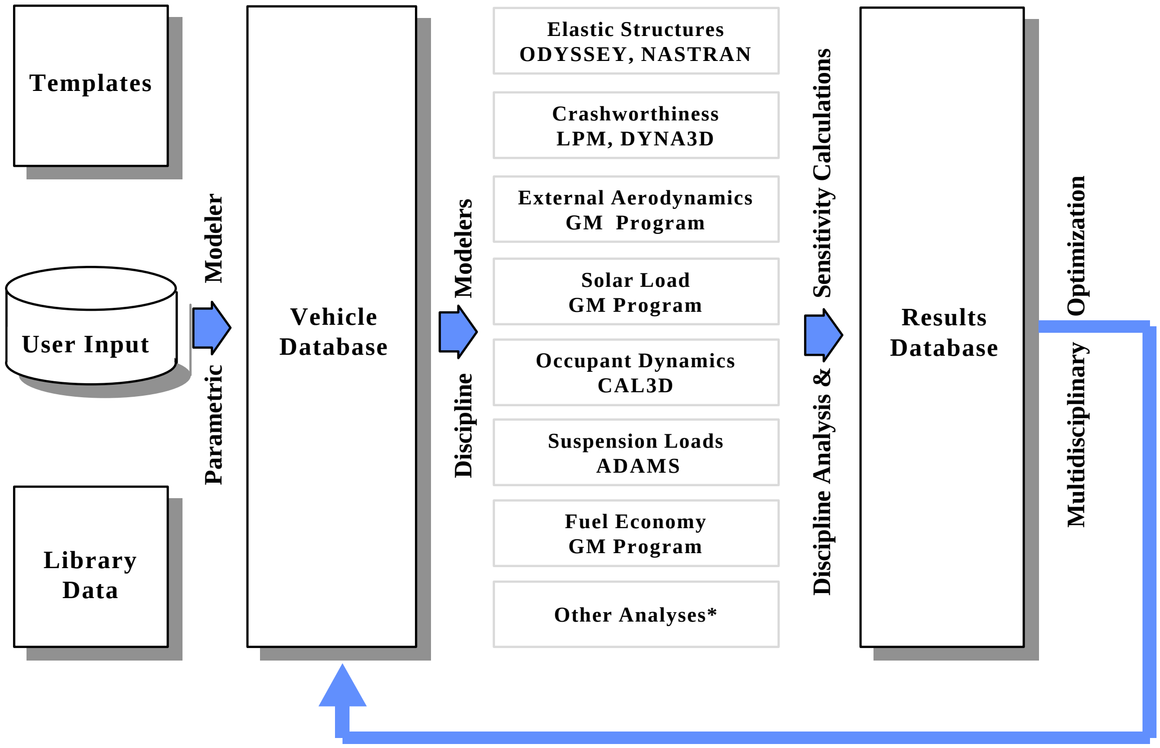

Figure 1. IVDA System Modules and Flow

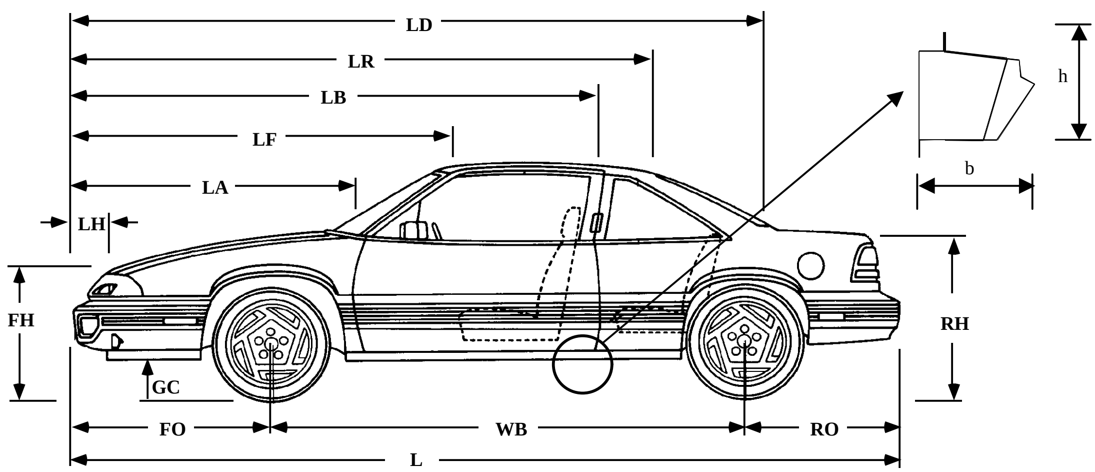

Figure 2. Overall Body Parameters and Typical Structural Parameters

As will be discussed later, an experimental multidisciplinary optimization capability has been added to the IVDA system. One of the goals of this system was to be able to examine several alternative ways of implementing the design strategy. For that reason there are currently few enforced sequences in the system and the control of the execution of the various modules are under the control of the various discipline and coordination human operators. To illustrate various aspects of this design process and its implementation in IVDA, we will begin by showing several examples. Each example is based on the simple problem outlined below but each uses a different design strategy. The final part of the report will discuss the technology used to implement these capabilities and some of its implications.

Examples

The example problem considered is to find a rear overhang (RO) that maximizes fuel economy. The available design variables are the total vehicle length (which in this parametric model expresses rear overhang since all quantities forward of the rear of the vehicle such as rear wheel location are not functions of the total length) and the traditional beam cross section sizing dimensions. The shape, other than lengthening, of the rear of the vehicle is not considered.

The underlying mechanics of the problem are that as the vehicle is lengthened, the drag will go down, which would tend to increase fuel economy. However, the mass of the vehicle increases which decreases the fuel economy. In addition, increasing the length of the vehicle decreases the fundamental bending and torsion frequencies. If these frequencies are below their target values, additional mass may be required for structural stiffening to bring the frequency back to its requirement.

In IVDA the structural analysis and design is handled by using a beam spring model and using the ODYSSEY/NASTRAN [2] programs to calculate response and gradients (sensitivities) and optimize mass for given constraints. An analysis takes approximately 10 minutes on a workstation and an optimization 1-3 hours. The aerodynamic drag is calculated by a neural net fit to test data so it is essentially instantaneous. Similarly, the fuel economy calculation is a rapid spread sheet calculation. Because both the aerodynamic and fuel economy calculations could have been replaced with more accurate and time consuming calculations, we will treat the process as if all three of the calculations required significant amounts of time. The flow of data is shown in Figure 3.

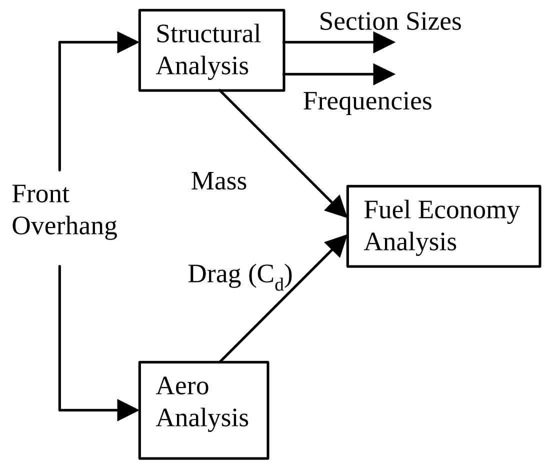

Figure 3. Information Flow in Example Problems

The traditional way to work this problem would be that every time information is required, a full cycle through the analysis is performed. That is full aerodynamic and structure calculations need to be made to calculate drag (C_d) and the structural mass must be adjusted to reflect any frequency requirements, then a fuel economy calculation can be made. Some design or optimization process would drive these calculations. Because this approach tends to require many calls to the analysis process, this tends to be a rather inefficient way to approach the solution.

Given the capabilities in a system such as IVDA, the above problem might be implemented in several more efficient ways, three of which we will illustrate. Each of these mimics a non-computer-based design strategy. In Example 1 the disciplines are only asked to provide local response information and after some point in the process are told to execute a design using only local design variables and local constraints. In Example 3, a roll down of requirements is initiated from the beginning, and much local design work is executed, but some iteration at the global level is required. Example 2 is an intermediate approach that makes use of local directional (sensitivity) information to guide the global design. This is often proposed as an initial step in flowing down requirements.

Table 1. No Initial Local Design Information

Step	Length mm	Fuel Economy Mpg	C_d	Mass Kg	Frequency (>25) Hz
0	4627	35.96	.292	1133.3	25.48
1	4677	36.04	.290	1134.8	25.44
2	4922*	36.17	.282	1142.3	24.65
3	4838*	36.15	.284	1139.7	25.04

Table 2. Local Design Information

Step	Length mm	Fuel Economy mpg	C_d	Mass Kg	Frequency (>25) Hz
0	4627	35.96	.292	1133.3	25.48
1	4677	36.04	.290	1134.8	25.44
2	4831*	36.18	.284	1138.0	24.95
3	4848*	36.16	.284	1139.2	24.97

Table 3. Roll Down of Design Available

Step	Length mm	Fuel Economy mpg	C_d	Mass Kg	Frequency (>25) Hz
0	4627	35.96	.292	1133.3	25.48
1	4677	36.04	.290	1134.8	25.44
2	4970*	36.15	.281	1144.8	25.00
3	4972*				

* Multidisciplinary Optimization Result

<u>Example 1: No initial local design</u>

The concept is that design variables naturally split into those that are of a global nature and those that are limited to a specific discipline. For this first strategy, only local response information is requested until the final step. This example is shown in Table 1. Step 0 represents a baseline design in which the structural cross section dimensions have been optimized for the given rear overhang (total length = 4627). The fuel economy, C_d, mass, and critical frequency constraint are also reported. Since only response information is returned, a step is required to generate directional (sensitivity) information. An arbitrary perturbation of 50mm is taken in Step 1. There is now sufficient information to generate a linear approximation with respect to length for all needed quantities; such as mass, C_d ,frequency and fuel economy. An optimization algorithm is then applied to identify an optimum length of 4922 which is labeled Step 2. This new length is then transmitted to the analysis codes and the values of C_d, mass, frequency, and fuel economy are calculated. Note that the frequency constraint of 25Hz is violated, primarily because the linear approximation was not sufficiently accurate. However, there now exists sufficient information to construct a quadratic approximation to all quantities. Using these approximations an optimization is again conducted, identifying an optimum length of 4838. When this length is returned to the analysis codes the remaining values in Step 3 are calculated. It would be possible to stop at this point, or the local analyses could conduct a local design in those variables that do not affect any of the other disciplines. In this problem this would be the cross-section design variables of the entire structure (there are 116 of these). This was done, requiring 3 additional structural analyses, and reduced the mass by

.2 kg which had no measurable effect on the fuel economy. This process initially used 4 analyses in each discipline plus 2 analyses and 2 sensitivity calculations in the final structural optimization.

Example 2: Local information available

This implementation uses more than just response information from the disciplines. The results for this approach are shown in Table 2. Step 0 is identical to Step 0 in Example 1. In the example problem, the structures disciplines can in fact generate sensitivities of the mass and frequency with respect to the cross-section design variables. The length variable is not directly known to the structures module so a sensitivity with respect to length cannot be generated. Therefore a Step 1 which is identical to the Step 1 in the first example must be made. Approximations with respect to the length (based on a linear response surface) and with respect to the section design variables (based on sensitivities for 116 variables) can now be made. The multidisciplinary optimization problem can then be solved using all 117 design variables. Only the length is shown in the table (4831). Upon re-analysis the frequency constraint was slightly violated (24.95). At this point a quadratic approximation based on length and updated sensitivity values for the section variables can be generated. The structural approximations based on length are not precisely correct because they contain now an evaluation in which the section variables as well as the length were changed. It is impractical to generate a response surface for all 117 variables since it would require a minimum of 118 analyses. Step 3 shows the results of the approximate multidisciplinary optimization (length 4848) and the subsequent full evaluations (frequency = 24.97, Fuel Economy = 36.16. This process used 4 analyses in each discipline and 2 structural sensitivity calculations. Note that no final local structural optimization was performed so there remains the possibility that this final design is not precisely optimum.

Example 3: Roll down of design

In this implementation advantage is taken of local design capabilities. The structural discipline capability has the ability to perform optimizations (designs) which minimize the mass subject to constraints. The results are shown in Table 3. Again Step 0 is the baseline and Step 1 is the 50 mm move, however the results shown in Step 1 are for a structurally optimized design in terms of the cross-section dimensions. Also the approximations used for the multidisciplinary optimization in Step 2 are based on optimized structural designs. The Step 2 multidisciplinary optimization produced a length of 4970. When the structural optimization was again performed the frequency constraint was initially infeasible (24.32), but the local optimization was able to resolve this and the final design from Step 2 showed a frequency value of 25Hz. Now quadratic approximations can be built based on these optimized results. The multidisciplinary optimization in Step 3 then produces a value essentially identical to the length for Step 2 so we can conclude that the design has converged. This process used 4 analyses in each discipline with 6 additional analyses and sensitivity calculations for structural optimization.

Discussion of Examples

Because of the nature of the problem solved it is not possible to draw firm conclusions about either vehicle design trends or the nature of which design strategy is best. What has been shown is a computer implementation that will allow these multidisciplinary problems to be handled mathematically and will allow different design strategies to be applied. The following section of the paper will discuss these issues. However, first, some observations based on the examples can be made.

The numerical differences among the quantities are in many cases extremely small. However, throughout the many exercises of these examples there has been sufficient consistency in the results to suggest that they are not being driven by numerical noise. All of the designs consistently allowed the length to increase, which means that the gain in fuel economy from decreased drag offsets the decrease in fuel economy due to the increase in mass. However, once the frequency constraint was encountered, the additional mass required to meet the constraint at the longer lengths eventually overrode the fuel economy gains due to the increased length. The range of final lengths (4838-4970) and Fuel Economy (36.15 - 36.16) suggest a rather flat optimum over a relative wide range of lengths. To reliably select a true optima from these designs is probably impossible with the available level of accuracy in the analyses.

Similarly it is impossible to identify a best process from this simple example problem. All of the processes work relatively well in terms of efficiency and quality of answer since the two sets of design variables (length and section dimensions) are fairly well independent for this problem. In addition all examples were started from a structural design that was quite good (optimal) for the initial total length so the effects of large changes in the structural cross section design variables was eliminated. However, it is possible to see some of the relative strengths and weaknesses. The no local design approach appears to have the most difficulty in following the frequency constraint, but initially requires the least effort from the local disciplines. It is easy to believe that in a more highly coupled problem, this could lead to inefficiencies. The second approach, which brings in local sensitivity information if available, could be considered the most efficient from the standpoint of structural analyses required, but because of the information that is to be shared, is the most complex to implement either informally or mathematically. Ultimately this process will produce a large, but potentially simple, design problem at the global level. The last process which used local design required the most local design effort (finite element analyses). This will occur when there is coupling between the disciplines that are neglected in the roll down process.

As indicated previously, in order to implement and automate such a system, several new capabilities are needed. The remainder of this report will discuss these in light of the system that was used for these examples.

A Multidisciplinary Design System

.

Parametric/design variable modeling

One of the fundamental issues for successfully implementing a computer-based multidisciplinary design strategy is that each discipline must be able to communicate with other disciplines and the decision making process with the same set of design variables or parameters. If different disciplines use different descriptions of the same quantity or geometric entity they have no way to communicate, particularly mathematically. This says that some common parametric description, or a mapping among different descriptions, must exist. This was a fundamental concept of IVDA and resulted in a significant amount of the development effort.

Since the early 90's the CAD vendors have been evolving such a parametric capability for the geometric representations that they create. Similarly, some CAE vendors, notably the finite element structural analysis vendors, have been evolving optimization capabilities based on parameterized design models. It is logical that a bi-directional coupling could be established. In general this may not be easy since many disciplines have evolved their own geometric preprocessors which parameterize the discipline model in ways that are different from the CAD models. We will assume the existence of such a common parametric system in what follows since we wish to focus on the multidisciplinary design issues, however to insure appropriate implementation, it may be necessary for the MDO community to be actively involved in the evolution of this technology.

Approximate Problems

There is a large amount of heuristic and research information suggesting that the way engineering design is efficiently conducted is that a limited amount of high quality, time consuming, expensive information is collected and a simple approximation of this information is constructed, either heuristically or mathematically. This simple model is exercised to identify an improved design and this new design is then evaluated using the high quality and expensive method. This process is certainly used in the heuristic and test method of design and the current analysis based methods. From the research standpoint it has been well established in the structural optimization area that this approach reduces the computational effort by at least an order of magnitude for moderately complex problems. We will propose that the ability to create and handle approximations based on more refined data is required.

In some disciplines highly accurate, extremely fast analyses may exist. From our standpoint these become highly accurate approximations for which no reference to a more accurate analysis need be made.

Multidisciplinary Design Strategy

Given the above assumption that a set of approximations will be available there are several pieces to this strategy. First there will need to be some process to operate on the approximations to identify the new and improved design. Since this will operate on the cheap-to-execute approximations, we will assume that any strategy, including exhaustive search, could be used. In practice, exhaustive search many prove inefficient and the process probably would be selected to take advantage of the nature of the approximations. The next level of the strategy is how the approximations are generated. The final piece of the strategy is how the approximations, the designs based on these approximations and the more detailed analysis are interwoven. The three example problems show alternative implementations at this level of strategy. These examples suggest that one would not want to impose a strategy a priori.

Clearly at the core is the concept of the approximate models built on information in what might be called the results database. Therefore, we will begin our discussion with how these approximations might be constructed and managed. We will develop the concept of the IVDA results database throughout this discussion, but it is essentially where all of the relevant information that comes from the discipline analyses is located. In the CAD environment, many product data manager (PDM) systems anticipate storing a pointer to files of completed analyses. However, to make use of this information in a design system, some sort of data extraction is required. In the IVDA system, we required each discipline to supply specific subsets of their output data to the results database. The form in the results database was structured to meet the requirements of the mathematically based design processes envisioned. Although we located all of the information in one database, there is no reason that the information could not have been distributed in several discipline databases. Part of this vision was that the results database would live through several different multidisciplinary designs as opposed to one single execution of a precisely stated problem. Thus a persistent idea was that old analyses might be reused to construct new approximations for newly posed problems. A conceptual layout of the results database is shown in Figure 4.

In the following sections we will discuss several types of data stored in the results database and their relationships.

Response Approximations

The type of an approximation that we are considering here is a calculation that takes fractions of a second to execute on whatever the current compute platform is. These would generally fall into two categories. The first is some sort of standard mathematical form that would be suitable for any discipline. Forms such as Taylor series, polynomial response surfaces, and neural nets fall into this category. The second category contains models developed for one specific discipline. These could be spreadsheets, lumped parameter models, simple discrete models, or specialized response surface or neural net models. It is anticipated that the generalized mathematical models would be generated from data that exists in the results database. The discipline specific models would be either generated or enhanced from the data in the results database. Thus we clearly need to provide capabilities to interrogate the results database to create and update these models.

In IVDA we generate response approximations with up to quadratic terms using two different methods. In the first approach, we use discipline generated responses and sensitivities at a single design point to create linear response approximations. The second approach uses only the discipline generated response values at a number of points in the design space to create up to second order response surface approximations. Both types of approximations are stored in the same relation in the results database.

Approximations Based on Discipline Response Sensitivities

We chose to make the creation of an approximation a decision at the multidisciplinary level as opposed to the discipline level. This is so the multidisciplinary design process would "understand" the approximations it had available. On the other hand, we predicated the results database on the idea that a decision to place data in it was made by the local discipline. This essentially placed the burden on the local discipline to warrant that the data was correct and might have some potential value. Thus it was necessary to create a location to store discipline supplied sensitivities prior to the decision to elevate them to approximations to be used in the multidisciplinary design process. While response data can be fairly compact (responses and associated design variables), sensitivity data can be rather extensive and to store this information for each returned response may be prohibitive. For that reason, in the IVDA results database only one set of sensitivities for each discipline is kept, the last one returned. In practice this has proved cumbersome because it requires that before approximations are to be constructed, the "correct" response is the last one loaded. A module has been implemented which on command transfers the sensitivities for a particular response from the sensitivity relation to the approximation relation in the results database.

Approximations Based on Response Surfaces

Most response surface generation processes assume that they are provided with a set of responses and their corresponding sets of design variables. They then use a prescribed algorithm to develop an interpolation scheme that fits these points in some best way. Most optimization or design methods that use response surfaces presume that for every new problem, a start from no information is made and that the information to develop the response surfaces is provided (perhaps n+1 vectors of the n design variables for which the responses must first be calculated). We wish to operate in a situation in which several designs have already been created (i.e. the results database is partially populated) and we wish to use as much of the already generated information as possible. We do recognize that at any given point in time there may be insufficient information and some additional analyses may need to be executed, but we wish to minimize the amount of this that must be done.

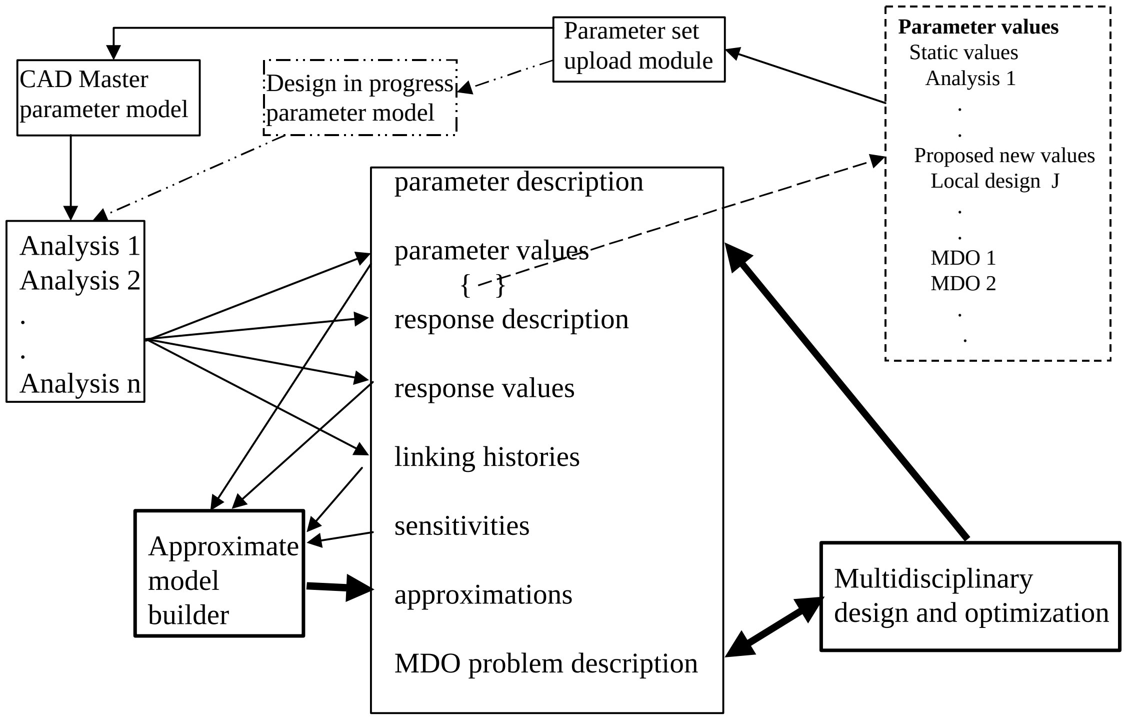

Figure 4. Results Database and Associated Modules

To accomplish this a module was developed which for a given response will identify in the results database all designs which calculated a value for this response. It then identifies the values of the design variables associated with these responses. At this point the information is available to submit to a standard response surface generator. There is a question as to whether all of the responses in the database are "good" in that some could have been loaded and later determined not to be valid. For that reason, it is possible for the user to specify which of the available responses are to be included in the fit.

The main difficulty with this process is associated with the potential interrelationships among the design variables. In a traditional regression approach, it is assumed that a unique set of design variables is identified and remains constant throughout the process. In this case the only concern is that the designs used for the response surface must not be linearly dependent in any fashion and there are standard methods to detect this situation. In the process proposed above these conditions cannot be guaranteed unless unreasonable restrictions are placed upon the design process.

Design Variable Linking with Approximations in Multidisciplinary Design

The fundamental concept in parametric modeling is that the number of degrees of design freedom are reduced by relating potential degrees of freedom to a reduced set of quantities by mathematical expressions. In the example problem described earlier the four points that describe the rear of the vehicle (upper and lower corners of the rear on each side of the vehicle) are all related to the total vehicle length. The process of constructing the relationships between potential design variables and a reduced set of actual variables for a given problem has been called linking in the optimization literature and that term will be used here. In order to allow for some amount of generality and future changes, we anticipate that it will be desirable to retain access to this extended set of potential design variables. Thus the basic set of design variables retained in the IVDA results database is not the set of design variables that are active on the current multidisciplinary design, but the complete set of design variables that is available in the template. For the examples shown earlier there are 3033 of these potential design variables. Again taking the examples described previously, the response surface module will identify that four of these variables (upper and lower, right and left) have been changed for any design that changes the rear overhang. However, for this problem, the local linking specifies that all four quantities move the same distance. Since the variables are not independent, we need to fit only one variable, not four. Therefore some process must be implemented that recognizes this situation and accounts for it. In the IVDA implementation this is accomplished by creating a historical link table relation. This relation contains the links that are used in any set of responses that are returned to the results database. It should be recognized that these links could have come from the central vehicle description, or they could have been provided and/or modified by the local disciplines. The response surface generating module then checks to see if the same linking was used for all of the responses to be used. If so, the number of design variables is reduced and the appropriate variables removed from the independent variable list for fitting. This allows a correct response surface approximation to be generated.

The difficulty here is that although several variables contributed to the total sensitivity calculation, all of the information is now attributed to one variable. This is fine as long as one wants to retain the current linking throughout the entire process. However, if information from another discipline did not contain this linking, and it was desired to allow linking changes in the multidisciplinary problem, difficulties could arise. Therefore it is desirable to decompose the linked, aggregated sensitivity into the components of the unlinked design variables. This can be done in an approximate sense by using the link relationships and a chain rule. Then based on the linking used in the multidisciplinary problem, the appropriate linked sensitivity can be reconstructed. For the specific example used here, the sensitivities would be equally split among the four points. This capability has been implemented. The difficulties with this approach are also obvious, since one might expect the sensitivities to be equal from side to side, but the two lower points could be expected to have different sensitivities than the two upper points.

In looking at the current state of parametric modeling implementation in CAD/CAE software, it is clear that many of these situations will arise here. It is recognized that the discipline parameterization must match the vehicle parameterization in terms of the quantities that are to be communicated. It is not however realized that the relationships among these quantities will change and that an interpretable record of these relationships may need to be kept. Just as it may be appropriate for a discipline to propose a change to a parametric dimension, a discipline may want to propose a change to the way these variables are linked, and the decision making process needs a history of these proposals.

Multidisciplinary Analysis with Approximations

Although we are dealing with approximations, the relationships among the various approximations for each discipline are the same as for the more complex modules. Therefore the issues associated with exchanging information are the same. There are two possible situations. Information from one discipline may flow forward. That is the output from one discipline may be input for the next discipline. For instance the loads calculated by a suspension program might be the input for a structural optimization program. Similarly there may be a feed back of information in which the output of one program is needed to calculate the input for another program whose output is the input for the first program. For instance the mass calculated by the structural optimization program is needed as input to the suspension program which calculates the load input for the structural optimization program.

If there is only feed forward, it is fairly easy to envision how a multidisciplinary design process would work: by properly ordering the analyses (or approximations), the outputs of the programs would be used as inputs for the following programs and all response properties could be properly calculated. An optimization capability could then be wrapped around the feed forward package.

If there is any feedback present, the process is more complicated since there will need to be inner loops around these feed back loops to insure convergence of the responses before data is passed to the next step in the process.

There are two approaches that might be implemented here. Since we are working with approximations that presumably execute very quickly, we could implement a full feedback and feed forward process that would express all of the interactions implied in the approximations. This in general will require developing approximations of the output quantities of each discipline module with respect to all of the input quantities, not just the design variables. For example in the example problem used here, an approximation of the fuel economy with respect to C_d and mass will be needed since these are the input quantities needed by the fuel economy module (Figure 3)

The second approach is through a mathematical formulation. Mathematically this situation can be expressed by what are called the global sensitivity equations.

$$\frac{dg_j}{dx_i} = \frac{\partial g_j}{\partial x_i} + \sum \frac{\partial g_j}{\partial g_k} \frac{dg_k}{dx_i} \tag{1}$$

In this first order approximation sense these equations allow us to create a complete approximation, dg_j/dx_i, to describe the effect of a design variable x on any output quantity g in terms of both feed forward and feed back. To do this we need the traditional derivatives, $\partial g_j/\partial x_i$, (approximations) of any output quantity with respect to its local input variables, plus the derivatives, $\partial g_j/\partial g_k$, (approximations) of any output response quantity with respect to any input response quantity. This is essentially a set of linear algebraic equations that can be solved by standard matrix methods.

Thus to implement either of the two approaches we need the same types of additional information, i.e. either approximations or sensitivities of response quantities with respect to input response quantities.

In the example problems, we had only a feed forward problem and response/response sensitivities were only needed for the fuel economy program. We used a first order approximation, calculating the sensitivities by finite differences. This was implemented in the first approach, treating the sensitivities as an approximation, and chaining the information through the approximations. This gives the same answer as the global sensitivity equations that in the case of feed forward reduce to a chain rule.

Multidisciplinary Design Strategies

As indicated earlier there are multiple levels to this strategy. We have proposed an approach in which the design or optimization strategy is applied to the approximations. While virtually any optimization package could be used, we used the feasible directions strategy in ADS [3] for the examples. Any of these packages require the availability of response and perhaps sensitivity information. Some process must be devised to interface the optimization algorithms with the approximation modules as described in the previous section. While we developed a simple input format that would point to the appropriate modules, one of the newer commercial MDO oriented packages could be adapted to the task. The output from the optimization package would then be a proposed new design that must be reloaded into the high level description for re-analysis, if accepted.

This brings us to the relationships between the high level vehicle description and the other parts of the process. Both IVDA and the commercial CAD packages envision a high level description of the present state of the vehicle which is the common central description which all analyses reference as their starting point. It is then envisioned that the local disciplines may explore alternative designs and propose a new set of design parameters. Most of the current CAD vendor thinking is around the process of allowing one discipline to upload its new set of parameters to the CAD model. It is unlikely, however, that all of a discipline’s proposed changes would be accepted. The more likely situation is that many disciplines may propose conflicting sets of design parameters and these conflicts must be resolved before a design can be uploaded to the central description. This is essentially the job that is handled by the multidisciplinary optimization process. Obviously, there needs to be some intermediate level of data storage to handle all of these proposed new designs. The results database in IVDA stores these proposed designs from all the disciplines. No discipline can directly input its results to the vehicle database. The only way that the vehicle database (CAD model) can be updated is through the results database. A module was created that will select a complete set of design parameters in the results database and return it to the vehicle database (Figure 4). This then treats the approximate multidisciplinary optimization as just another discipline that has returned a proposed design that can be selected for return to the results database.

Although it was not implemented in the current version of IVDA, it is reasonable to assume that an intermediate copy of the vehicle database will be needed to hold modifications of the design parameters that are used to construct approximations required by the design process, such as those required by the length variable in the structural discipline in the examples. This is shown in Figure 4 with broken lines.

The remaining issue is how a high level multidisciplinary design strategy will interact with the approximate design strategy and the complex analyses. In the example problems this was handled by direct interaction, implementing all of the necessary executions through high level IVDA commands that provide for the input generation, execution, and transfer of the results of the various disciplines. This makes this process more time consuming than necessary, but because the exact series of steps is as yet undetermined, it seemed inappropriate to automate them, until such time as the rest of the system is more formalized. It does suggest that a menu of appropriate actions should be generated to guide the user through the process.

Summary

By first using some simple example problems we have tried to motivate a view of a mathematically based design process that has parallels in the traditional processes. This vision involves an interplay between complex, time consuming analyses and approximations based on these analyses. It is unlikely that a single multidisciplinary design strategy will suffice for all problems. Therefore a system must evolve that will handle a number of different strategies. Three classes of issues were discussed. First, a method to share a common description of design parameters must be implemented, Associated with this is a necessity to keep track of the linking relationships among the potential set of design variables as these may change throughout the design process. Second is the need to have for each discipline a quickly executed approximation of a perhaps more complex behavior. These approximations can either be supplied by the disciplines, for example sensitivities if available, or they might be created at a higher level by examining all of the available detailed analysis results. The IVDA system was constructed explicitly to examine the latter situation and required additional sophistication to implement. Finally, if the previous two capabilities are in place, a shared set of design parameters and a shared set of approximations, the implementation of a design or optimization strategy is fairly straightforward and a wide range of strategies can be implemented including heuristic and mathematically based strategies.

Acknowledgements

The authors would like to acknowledge the valued contributions of V. Sankar who did much of the development programming of the multidisciplinary capability described here and who contributed to the discussions of the basic strategy. We would also like to acknowledge Bob Lust for pointing out the use of the chain rule and the linking information to develop a more generalized approximation to the decoupling of the aggregated sensitivity information.

References

1. J. Bennett, et. al., “A Multidisciplinary Framework for Preliminary Vehicle Analysis and Design,” GM Research Publication R&D-8290, Feb 9,1995.

2. M. E. Botkin, et al, “Structural Sizing Optimization Using an External Finite Element Program,” Proceedings of 28th Structures, Structural Dynamics and Materials Conference, AIAA, New York, 1987.

3. G. N. Vanderplaats, H. Sugimoto, and C. M. Sprague, ADS-1: A New General-Purpose Optimization Program,” Proceedings of 24th Structures, Structural Dynamics and Materials Conference, AIAA, New York, 1983.